UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/603,251	06/25/2003	Thomas J. Heller JR.	POU920020123US1	3683

7590 03/22/2007

Lynn L. Augspurger
IBM Corporation
2455 South Road, P386
Poughkeepsie, NY 12601

EXAMINER	
PARIKH, KALPIT	
ART UNIT	PAPER NUMBER
2187	

SHORTENED STATUTORY PERIOD OF RESPONSE	MAIL DATE	DELIVERY MODE
3 MONTHS	03/22/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.

Office Action Summary	Application No.	Applicant(s)	
	10/603,251	HELLER ET AL.	
	Examiner	Art Unit	
	Kalpit Parikh	2187	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *23 June 2003*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-15* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-15* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *23 June 2003* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

The instant application having Application No. 10/603251 has a total of 15 claims pending in the application; there are 1 independent claim and 14 dependent claims, all of which are ready for examination by the examiner.

I. **INFORMATION CONCERNING OATH/DECLARATION**

Oath/Declaration

The applicant's oath/declaration has been reviewed by the examiner and is found to conform to the requirements prescribed in **37 C.F.R. ' 1.63.**

II. REJECTIONS NOT BASED ON PRIOR ART

Claim Objections

Claim 13 objected to because of the following informalities: Claim 13 recites "a processors storage transaction." It appears the claim was intended to recite "a processor's storage transaction." Appropriate correction is required.

Claim Rejections - 35 USC § 112

1. The following is a quotation of the second paragraph of 35 U.S.C. 112:

   The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

2. **CLAIMS 1-15** are generally narrative and indefinite, failing to conform with current U.S. practice. They are replete with grammatical and idiomatic errors.

3. **CLAIMS 1-15** rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

   Claim 1 recites "coherent cache regions in the system are required to examine a coherency transaction." It is unclear how a region of cache is capable of performing the act of examining. Particularly, the claim does not recite what constitutes a coherent cache region, and therefore a coherent cache region is understood to mean a region of memory.

   Claim 2 recites the limitation "the hardware." There appears to be insufficient antecedent basis for this limitation in the claim.

   Claim 2 recites "the total processors." There appears to be insufficient antecedent basis for this limitation in the claim.

Claim 2 recites "a node of the computer." It is unclear if "a node of the computer" is refererring to one of the plurality of nodes.

Claim 3 recites the term "large." The term "large" a relative term, which renders the claim indefinite.

Claim 3 recites "multiple instances of a node." It is unclear how there can be multiple instances of a node.

Claim 5 recites the limitation "the mode bits." There appears to be insufficient antecedent basis for this limitation in the claim.

Claim 6 recites "provided by provided by supervisor software or firmware." It is unclear what is meant by "provided by provided."

Claim 6 and 7 recite "logical partitions are provided." It is unclear what "the logical partitions" are of.

Claim 7 recites "each partition." There appears to be insufficient antecedent basis for this limitation in the claim.

Claim 7 recites "the total number of processors." There appears to be insufficient antecedent basis for this limitation in the claim.

Claim 7 recites "

Claim 8 and 10 rejected under 35 U.S.C. 112, second paragraph, as being indefinite in that it fails to point out what is included or excluded by the claim language. The claims appear to be reciting the benefit of the invention (i.e., "cache coherency is optimized").

Claim 9 recites "the hardware." There appears to be insufficient antecedent basis for this limitation in the claim.

Claim 11 recites "The physical DRAM storage." There appears to be insufficient antecedent basis for this limitation in the claim.

The term "large" in claim 12 is a relative term, which renders the claim indefinite. The term "large" is not defined by the claim, the specification does not provide a standard for ascertaining the requisite degree, and one of ordinary skill in the art would not be reasonably apprised of the scope of the invention.

Claim 15 rejected under 35 U.S.C. 112, second paragraph, as being indefinite in that it fails to point out what is included or excluded by the claim language. The claim appears to be reciting the benefit of the invention.

III. REJECTIONS BASED ON PRIOR ART

Claim Rejections - 35 USC § 102

5. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.
> (e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

6. **CLAIMS 1-6,8-11 AND 13-15** rejected under 35 U.S.C. 102(e) as being anticipated by Morioka et al. (US Pat No. 6631447).

As per claim 1, Morioka et al. disclose a multiprocessor computer system comprising: a computer system (see FIG 1: 100) having

- a plurality of processing nodes (see FIG 1: 200) and caches (see FIG 1: 400) and a node controller (see e.g., FIG 1: 300)
    - o which use processor state information (FIG 7 AND FIG 2) according to mappings provided by supervisor software or firmware (see e.g., COL 14 58-62) of allowable physical processors (see e.g., FIG 7: 3301) to an application workload (see e.g., FIG 7: "TASK 0")
        - ▪ to determine which coherent cache regions in the system are required to examine a coherency transaction produced by a single originating, processor's storage request (see e.g., FIG 6: "CACHE COHERENCY AREA").

As per claim 2, Morioka et al. disclose

> The multiprocessor computer system according to claim 1 wherein a node of the computer **(see e.g., FIG 14: 100)** has dynamic coherency boundaries **(see e.g., FIG 14: LCC AND GCC)** such that the hardware uses only a subset of the total processors in the system for a single workload at any specific point in time **(see e.g., FIG 7)** and optimizes the cache coherency as the supervisor software or firmware expands and contracts the number of processors which are being used to run any single workload **(see e.g., ABSTRACT).**

[Scope of this limitation appears to be indeterminate because the claim recites the hardware optimizes coherency but does not further recite how to perform said optimization.

Morioka et al. appears to teach the limitations of the claim because Morioka et al. teaches that as a processor executes transactions of a task, the coherency check for said transactions is preformed over either a local region or a global region based on the LCC/GCC designation of the transaction.]

As per claim 3,

> The multiprocessor computer system according to claim 1 wherein multiple instances of a node are connected with a second level controller to create a large multiprocessor system **(see e.g., FIG 15: 200).**

As per claim 4,

> The multiprocessor computer system according to claim 1 wherein said node controller uses mode bits **(see e.g., FIG 5: LCC AND GCC)** to determine which processors must receive any given transaction that is received by the node controller **(see e.g., COL 12 LINES 24-63 AND FIG 5).**

As per claim 5,

> The multiprocessor Computer system according to claim 1 wherein a second level controller is provided which uses the mode bits to determine which nodes must receive any given transaction that is received by the second level controller **(see e.g., FIG 14: 1800)**.

As per claim 6,

The multiprocessor computer system according to claim 1 wherein logical partitions are provided **(see e.g., FIG 2: "PHYSICAL ADDRESS SPACE")** and mapping of logical partitions to allowable physical processors **(see e.g., FIG 7)** is provided by provided by supervisor software or firmware of allowable physical processors to an application workload **(see e.g., FIG 7: TASK #0 = CLUSTER #0 AND CLUSTER #1)**.

[FIG 7 appears to disclose a mapping between partitions to physical processors to an application workload (i.e., "TASK").]

As per claim 8,

The multiprocessor computer system according to claim 1 wherein a single workload uses only a subset of the total processors in the computer system for a single workload at any specific point in time **(see e.g., FIG 7: TASK # 0)** for an assigned partition **(see e.g., FIG 7: TASK #0 VIRTUAL ADDRESS SPACE)** and a distinct cache coherency is optimized for the address space of the assigned partition as the supervisor software or firmware expands and contracts the number of processors which are being used to run any single workload in said assigned partition

[The claim recites cache coherency is optimized for the address space but does not further recite steps to perform such action.

The system of Morioka et al. is understood to optimize cache coherency region as applications are run on the system.]

As per claim 9,

The multiprocessor computer system according to claim 1 wherein a single workload uses only a subset of the total processors in the computer system for a single workload at any specific point in time **(see e.g., FIG 7: TASK #0)**, and multiple cache coherent regions are assigned for different partitions as more independent workloads coexist on the same hardware

[Each local cache coherent region of a cluster is considered to be a different cache coherent region. Morioka et al. discloses independent workloads can be assigned to the hardware (see e.g., FIG 7: TASK #0 and TASK #1).]

As per claim 10,

The multiprocessor computer system according to claim 1 wherein cache coherence regions encompass subsets of the total number of processors **(see e.g., FIG 7 AND FIG**

14: LCC CLUSTER 0) and caches in the computer system and a single workload uses only a subset of the total processors in the computer system for a single workload at any specific point in time **(see e.g., FIG 7: TASK #0)** for an assigned partition **(see e.g., FIG 7: TASK #0 VIRTUAL ADDRESS SPACE)** and a distinct cache coherency is optimized for the address space of the assigned partition as the supervisor software or firmware expands and contracts the number of processors which are being used to run any single workload in said assigned partition.

[The claim recites cache coherency is optimized for the address space but does not further recite steps to perform such action.

The system of Morioka et al. is understood to optimize cache coherency region as more programs are run on the system.]

As per claim 11,

The multiprocessor computer system according to claim 1 wherein software and/or firmware define which subset of processors in a large multiprocessor must participate in a coherency transaction independent of which processing node is connected to the physical DRAM storage being requested by said single originating processor **(see e.g., FIG 9).**

[The determination to decide which cache coherency regions must participate in cache coherency is independent of which processing node is connected to the physical DRAM.]

As per claim 13,

The multiprocessor computer system according to claim 1 wherein cache coherence mode bits are appended to a processors storage transactions when transmitted to a connected processor of said multiprocessor computer system **(see e.g., FIG 5 "LCC/GCC-REAL ADDRESS")**.

As per claim 14,

The multiprocessor computer system according to claim 13 wherein said cache coherence mode bits are used in a decision determining whether the single originating processor's storage request must be transmitted to additional processors in the system **(see e.g., FIG 11: 1201).**

As per claim 15,

The multiprocessor computer system according to claim 14 wherein an increase in the effective utilization of the address bandwidth of the buses used to interconnect the processors of a multiprocessor system allows movement of workload among physical

processors in a multiprocessor system at the same time as a reduction of the address bandwidth required to maintain cache coherency among all the processors.

[The claim appears to be reciting the benefit of the invention (see Specification PAGE 8 LINES 3-10). Such benefits are understood to be present in the system of Morioka et al. (see Morioka et al. (ABSTRACT).]

7. **CLAIMS 1,6 AND 11-12** rejected under 35 U.S.C. 102(b) as being anticipated by Hagersten et al. (US Pat No. 6226671).

As per claim 1,

A multiprocessor computer system comprising, a computer system having a plurality of processing nodes **(FIG 2: 30A AND 30B)** and caches **(COL 1 LIENS 35-38)** and a node controller **(FIG 2: 14A AND 14B)** which use processor state information according to mappings provided by supervisor software or firmware **(see e.g., FIG 4)** of allowable physical processors to an application workload **(see FIG 3 AND FIG 4 AND COL 7 LINES 63-67)** to determine which coherent cache regions in the system are required to examine a coherency transaction produced by a single originating processor's storage request **(see FIG 6).**

As per claim 6,

The multiprocessor computer system according to claim 1 wherein logical partitions are provided **(see e.g., FIG 2: PHYSICAL MEMORY)** and mapping of logical partitions to allowable physical processors is provided **(see e.g., FIG 2: 130, 132, AND 134)** by provided by supervisor software or firmware of allowable physical processors to an application workload **(see e.g., COL 7 LINES 63-67).**

As per claim 11,

The multiprocessor computer system according to claim 1 wherein software and/or firmware define which subset of processors in a large multiprocessor must participate in a coherency transaction independent of which processing node is connected to the physical DRAM storage being requested by said single originating processor **(see e.g., FIG 6).**

[The determination to decide which cache coherency regions must participate in cache coherency is independent of which processing node is connected to the physical DRAM (see COL 6 LINES 64-COL 7 LINES 5).]

As per claim 12,

The multiprocessor computer system according to claim ii wherein the movement of a process between nodes of a large multiprocessor is effectuated without moving physical storage contents and without requiring subsequent broadcasting of the storage references originated by the process from said single originating storage request to all of the caches in the multiprocessor (see e.g., COL 8 LINES 21-31).

[Hagersten et al. teaches that while initially the storage references are treated global after a process has migrated to a remote node, subsequent processing may cause the storage reference to be treated local. Subsequent storage references would not require global broadcast (see e.g., COL 8 LINES 58-61).

Further the claim does not appear to recite which physical storage contents are moved.]

Claim Rejections - 35 USC § 103

8. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

9. **CLAIMS 7** rejected under 35 U.S.C. 103(a) as being unpatentable over Hagersten et al. (US Pat No. 6631447) in view of IBM (NPL: 2002).

As per claim 7, Hagersten et al. disclose the multiprocessor computer system according to claim 1 wherein logical partitions are provided (see e.g., Hagersten et al.: FIG 2) for the supervisor software or firmware (see e.g., Hagersten et al.: FIG 4) which maps allowable physical processors (see e.g., Hagersten et al.: FIG 2: 130,132,134) to an application workload (see e.g., Hagersten et al.: COL 7 LINES 63-67) and the supervisor software or firmware assigns cache coherence regions which encompass subsets of the total number of processors and caches in the system chosen for their physical proximity and defines a distinct cache coherency region for each partition (see e.g., .

However, Hagersten et al. do not expressly disclose a hypervisor assigns cache coherence regions.

In the same field of endeavor IBM (NPL) discloses a hypervisor can be utilized to create a secure partition environment (i.e., cache coherency region in the system of Hagersten et al.) for operating systems through logical partitioning (see e.g., IBM (NPL): PAGE 2 FIG AND PAGE 4).

At the time of the invention it would have been obvious to one of ordinary skill in the art to modify the system of Hagersten et al. to utilize a hypervisor to perform the partitioning of physical memory in order to create the cache coherency regions of Hagersten et al. (see Hagersten et al.: FIG 2).

The suggestion/motivation for doing so would have been to increase partition integrity though the use of a hypervisor (see IBM (NPL): PAGE 4 LEFT COLUMN).

Therefore it would have been obvious modify the system of Hagersten et al. by using a hypervisor, as taught by IBM, to partition the physical memory into logical partitions (which are treated as cache coherent regions in the system of Hagersten et al.,) for the benefit of increased partition integrity to arrive at the invention as specified in the claims.

Further it is noted that Applicants' disclosure makes clear that a hypervisor performing the functions recited in the claim is known in the art (see PAGE 2 LINE 32- PAGE 3 LINE 12).

IV. RELEVANT ART CITED BY THE EXAMINER

The following prior art made of record and not relied upon is cited to establish the level of skill in the applicant's art and those arts considered reasonably pertinent to applicant's disclosure. See **MPEP 707.05(c)**.

The following references teach a selective snooping operation.

U.S.PATENT NUMBER	FIGURES
5860114	FIG 4
7003633	ABSTRACT AND FIG 2

V. CLOSING COMMENTS

Conclusion

Va. STATUS OF CLAIMS IN THE APPLICATION

The following is a summary of the treatment and status of all claims in the application as recommended by **M.P.E.P. ' 707.07(i)**:

Vb. CLAIMS REJECTED IN THE APPLICATION

Per the instant office action, claims 1-15 have received a first action on the merits and are subject of a first action non-final.

For at least the above reasons it is the examiner's position that the applicant's claims are not in condition for allowance.

VI. DIRECTION OF FUTURE CORRESPONDENCES

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Kalpit Parikh whose telephone number is (571) 270-1173. The examiner can normally be reached on MON THROUGH FRI 7:30 TO 5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Donald Sparks can be reached on (571) 272-4201. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Kalpit Parikh/

Kalpit Parikh
Examiner
Art Unit 2187

March 12, 2007

DONALD SPARKS
SUPERVISORY PATENT EXAMINER